IAC-17.A1.6.1x40187

FORWARD CONTAMINATION OF OCEAN WORLDS: A STAKEHOLDER CONVERSATION

Brent Sherwood
Jet Propulsion Laboratory, California Institute of Technology, Pasadena CA USA, brent.sherwood@jpl.nasa.gov

Adrian Ponce
Jet Propulsion Laboratory, California Institute of Technology, Pasadena CA USA, Adrian.Ponce@jpl.nasa.gov

Michael Waltemathe
Ruhr-Universität Bochum, Bochum, Germany, michael.waltemathe@googlemail.com

Abstract

A fundamental requirement for space missions designed to touch "potential habitats" is the single number 10^{-4}, the allowable probability of a single Earth organism contaminating the potential habitat. Many aspects of a mission that affect its complexity and cost – hardware design and manufacture, assembly and test, and mission operations – are driven by this value, so it is important, on the threshold of an era of exploring ocean worlds, to have confidence in it. Yet despite its long pedigree and occasional reviews, we find that the current requirement lacks programmatically defensible justification. At issue are three weaknesses: 1) microbial biology, in particular the science of extremophiles, is a rapidly changing field; 2) forward contamination is both a scientific and an ethical issue, yet no ethics-based conversation is apparent within policy-setting circles; 3) because of these two factors, policy-setting cannot be static. We review the history of the requirement; how the evolving understanding of biology could drive it up or down; how the forward-contamination hazard relates to risk-management practice and to the ethics profession; and how a contemporary stakeholder conversation could adapt lessons already learned by other fields.

Keywords: Planetary Protection, Forward Contamination, Ocean Worlds, Microbial Biology, COSPAR, Ethics

Forward Contamination: What it Is and Why We Care

A conversation about forward contamination is becoming more urgent.

Forward contamination means the inadvertent introduction of viable Earth life into an offworld environment that could support it. For human space flight this act – introducing life into space and sustaining it there – is of the essence. Humans, after all, depend on rich microbiomes that need to be brought along, even if thoughtfully.

But for humanity's exploration vanguard – machine avatars pursuing direct evidence of potential life in niches among the many ocean worlds of our solar system – forward contamination poses a significant, systems-level obligation and opportunity for technological excellence.

Scientific and public interest in ocean-world exploration is intensifying. For example, awareness is building that Mars, which has always fascinated humankind, is the closest and largest, but driest, example of something even more extraordinary: a whole class of ocean worlds, spread across the expanse of the solar system (Figure 1) [1].

These worlds are diverse, intriguing, and accessible to us in this century if we lean into it. Some of them have vast interior oceans of salt water, dwarfing Earth's surface seas, and having both a seafloor and a 'seaceiling'. We already have evidence of propitious conditions including warm hydrothermal activity, in one of them, Enceladus (Figure 2). Questions abound: what is the chemistry of such oceans…are some of them habitable? …and is there evidence of life in any of them? …more than one of them? …how does it work? Here we can learn the limits of life in the cosmos [2].

The dozen or more ocean worlds available to us offer manifold research opportunities in several fields that can be driven by the opportunity to make exobiology a flight science. Many expeditions would be needed, each carrying the best space flight technologies we can imagine. Together they would conduct an unprecedented form of biological and oceanographic exploration. Such an OWEP (ocean worlds exploration program) could be one of the 21st century's grand scientific adventures: able to encourage nations to collaborate, focusing their respective talent and innovation for many decades. Through such an exploratory adventure, we would directly inventory the presence or absence of life throughout the tiny fraction of the universe that is physically accessible to humankind. Astronomy is beginning to reveal startling diversity among billions of exoplanetary systems. Yet in only twelve or so places – the ocean worlds of our solar system – can we ever tangibly contact non-Earth life, should it exist. This constitutes at once a sobering limit on eventual human knowledge in a limitless universe, and a perishable

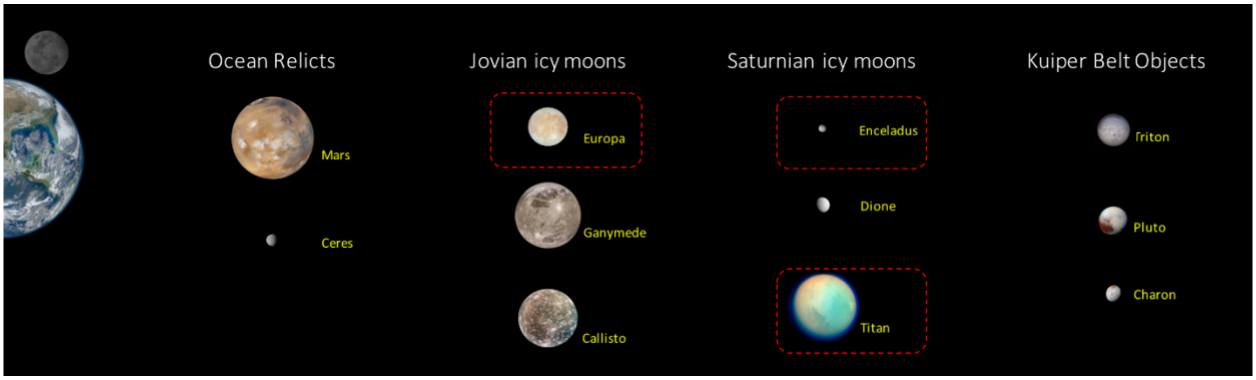

Figure 1. About a dozen ocean worlds are known in our solar system. Most have interior oceans, which contain the entire inventory of physical evidence for life off Earth that humanity will ever have. Europa, Enceladus, and Titan are the most key.

resource since the act of exploring it could compromise it for all future generations.

We know where to look and we are developing the means to do it. Instruments and mission systems up to the task of exploring icy moons in deep space are budding into maturity. Answers to age-old questions are coming within our grasp. It is even conceivable that inventory could be completed within this century – making now an epic time for humankind.

Multiple mission projects are already in development, with more in formulation and proposal pipelines. Now is a near-optimal time to consider afresh the forward-contamination (FC) requirement: one of the fundamental rules governing all such investigations. Dozens of potential missions in this century would depend on it.

Most especially, the generations that would implement these missions over the decades, and manage their consequences, deserve the opportunity to infuse their ideas and priorities into planning such a momentous quest. Such a conversation would be multi-faceted:

1. Humankind's potential to contaminate the ocean worlds (technical probabilities)
2. The implications of doing so (science-based and ethics-based discussions and decisions)
3. Acceptable ways of managing this risk (program planning and execution)

…and would therefore require collaboration by a diverse community of technical and non-technical stakeholders.

That type of collaboration, in addition to simply shaping the smartest possible program, could become a potent influence for sustaining an OWEP over time. A modest annual program investment, estimated to cost about 1/40th more than today's total NASA budget (Sherwood et al., 2017), may be able to inventory the life-bearing potential of our ocean worlds. Anything more, such as aggressive simultaneous exploration of multiple potential habitats across the solar system, is conceivable but most likely would depend on discovering promising leads. For example, the remarkable growth and momentum of today's Mars Exploration Program (MEP) followed tantalizing and later, rewarding, findings: putative biosignatures in a Mars meteorite found on Earth, then a campaign to "follow the water," culminating today in ample evidence of ancient, long-term clement conditions and clear ideas of what to do next and where to do it. Without this sequence of promising results, scientific exploration of Mars would likely have been far slower.

Today's FC requirement is simple and clear [3]:

> *Limit to 10^{-4} the probability that a single viable Earth cell is introduced into a potential habitat, defined as liquid water or warm ice.*

It originated in the 1960s, as we were preparing to land on Mars for the first time with Viking, and was adopted by agreement at that time because it appeared to balance achievability with the goal of protecting science.

The requirement has been formally reviewed several times since its origin; no review has determined a sound basis for changing it. But the requirement's context continues to change significantly. We are now confronted by a range of distinct worlds rich in water, armed with affordable capabilities to start exploring them, and informed by rapidly evolving knowledge about life – what it is, how it works, where it can be found on Earth, what it needs to be viable.

Eventually, returning samples from these places would require that the world be ready, technically and societally, for receiving them. The forward-contamination requirement allows us to engage with many of the same stakeholders, and demonstrate many of the same technologies, first on a problem without any physical risk to humanity. And the society that could mount such missions is changing altogether around us as we plan missions: people's interests, priorities, and beliefs are highly fluid. All of humanity will share in both the thrill of new existential knowledge and the consequences of how it is gained.

This analysis summarizes the history of how the requirement became what it is; articulates a rationale for its reconsideration; and offers cautions about application of risk analysis techniques. Importantly, it explores the ethical dimension that has been neglected since the original debates in the 1960s and 1970s. Finally, it suggests a method and criteria for catalyzing and conducting a broad stakeholder conversation so that mission planners can move forward confidently into this new frontier.

History of the Current Requirement

Exploring ocean worlds in the outer solar system would significantly advance our opportunity to answer the question "are we alone?" Yet it would also simultaneously advance a key scientific risk: that microbial contamination might destroy forever our ability to answer the question.

This challenge of protecting a potential extraterrestrial biosphere is not new. It was first identified in 1956, a year before Sputnik was launched, by Nobel laureates Joshua Lederberg and Melvin Calvin, who stated "that the potential for scientific discovery could be forever compromised if space exploration was conducted without heed to protecting the environments being explored" [4]. These warnings led the US National Academy of Sciences (NAS) to call upon the International Council of Scientific Unions (ICSU, now International Council for Science) for an international response. In 1958, ICSU first established the Committee on Contamination by Extraterrestrial Exploration (CETEX), which met only twice; and then the Committee on Space Research (COSPAR), which has played the keystone role in international development of planetary protection policies.

Also in 1958, the NAS established the Space Science Board (SSB, later renamed Space Studies Board), which became the major adviser to NASA (also formed later in 1958) on all interplanetary contamination issues. In 1961, ICSU declared that all countries launching space experiments that could have an adverse effect on other scientific research should provide ICSU and COSPAR with the information necessary to evaluate potential contamination. From the beginning, a major objective of COSPAR was to open a dialog between Eastern and Western bloc scientists [4]. In 1962, USSR Chairman Khrushchev wrote to US President Kennedy regarding "heavenly matters," in which he urged that "any experiments in outer space which may hinder the exploration of space by other countries" should be discussed, and agreements reached "on a proper international basis" [4].

A 1963 JPL engineering study led by L.D. Jaffe assessed microbial contamination probabilities for Mars exploration [5]. The high sterility assurance levels needed to protect biologically interesting extraterrestrial environments require approaches – such as dry-heat baking, irradiation, and ethylene oxide gas bath – originally developed for the food cannery industry in the 1940s [6]. However, appropriate microbial reduction parameters for planetary protection rest on a narrow parameter space that "threads the needle" between two unacceptable results: failure of sensitive electrical equipment; and survival of microorganisms [4].

Ultimately, the Jaffe study concluded that a 10^{-4} probability of contamination was appropriately balanced: it would keep the chance of contamination low compared to the chance of other causes preventing useful biological data from Mars from being obtained [5]. Specifically, the study first calculated the probability to be $10^{-3.5}$ that no useful biological information would be returned over a multi-mission program, based on the assumption that 75% of 28 flights over 14 trajectory opportunities would fail to return information (50% might fail to reach Mars, and 50% of those that would reach Mars might fail to return information). This reasoning led to rationalizing that the probability of the whole program contaminating Mars should be held lower than $10^{-3.5}$.

The Jaffe study also considered a second rationale: that "the chance of contaminating Mars in the course of unmanned exploration should be kept low compared to

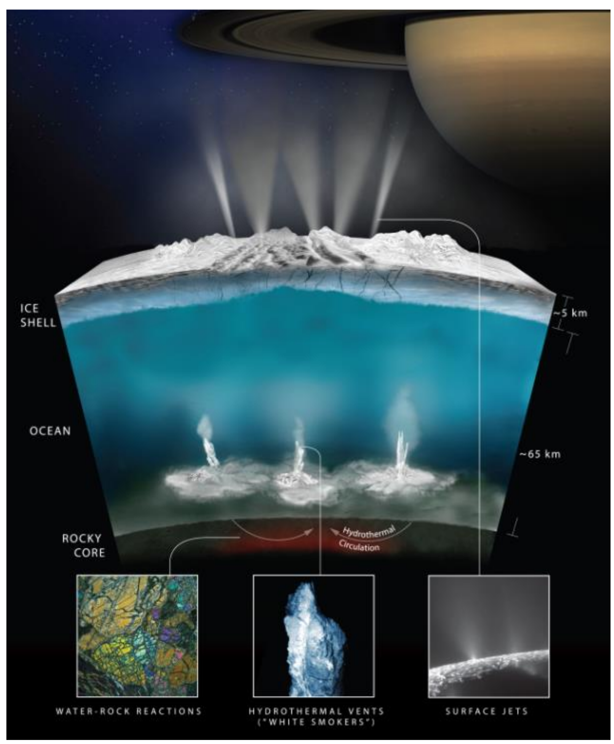

Figure 2. Cassini discovered that Enceladus has a global interior ocean, alkaline salt water supporting organic chemistry, and with warm hydrothermal systems active in its seafloor. A continuous plume expresses the ocean material directly into space where it can be sampled.

the chance of contaminating it the first time a manned landing occurs". It assumed that human exploration of Mars would have a 10% probability of contamination, which would make 10^{-2} an appropriate limit for probability of contamination by robotic precursors. The study concluded that contamination probabilities of $10^{-3.5}$ and 10^{-2} per mission from these two considerations, respectively, would yield $10^{-4.6}$ and $10^{-3.1}$ for 14 successful missions out of 28, so that "perhaps an intermediate value of about 10^{-4} is reasonable."

A conference on spacecraft sterilization sponsored by the NASA Biosciences Programs in 1962 focused on the Jaffe study, and adopted the 10^{-4} probability limit as a desirable goal. The 10^{-4} value was also suggested by an SSB study [7] for flyby missions, as an alternative to sterilization.

These guidelines, however, did not mandate a universal standard of planetary protection, and COSPAR, which governs international space policy, conducted its own deliberations on acceptable contamination levels. In 1964, COSPAR established a probabilistic framework for developing planetary protection standards, advocating "a sterilization level such that the probability of a single viable organism aboard any spacecraft intended for planetary landing or atmospheric penetration would be less than 10^{-4}, and a probability limit for accidental planetary impact by unsterilized flyby or orbiting spacecraft of 3×10^{-5} or less…during the interval terminating at the end of the initial period of planetary exploration by landing vehicles" [4]. COSPAR further resolved that the probability of a planet of biological interest being contaminated within the period of biological exploration should be no more than 10^{-3}, and that this standard should be adopted by all states engaging in the exploration of space [4].

In 1966, COSPAR suballocated the 10^{-3} overall limit to 4.4×10^{-4} each for the US and USSR, and to 1.2×10^{-4} for other spacefaring nations. For the case of the Viking landers, the integrated probability assignments each received initial suballocations of 7.2×10^{-5}, but because of the successful Mariner Mars missions this was augmented to a suballocation of 10^{-4} per Viking lander [4]. In 1967, NASA issued directive NPD 8020.10 to harmonize with these COSPAR requirements.

In 1967, the US, USSR, and UK strengthened their commitment to, and leadership on, planetary protection by negotiating, and becoming states party to the Outer Space Treaty [8]. Having signed the treaty binds the US under both international law and the US Constitution to "pursue studies of outer space, including the moon and other celestial bodies…so as to avoid their harmful contamination" [9]. To abide by the treaty's imperatives, NASA adhered to SSB recommendations to formulate a planetary protection policy for submittal to COSPAR for approval. The Outer Space Treaty has since been signed and ratified by 104 nation states; another 24 have signed but not ratified.

After the 1977 Viking missions revealed the surface of Mars to be more inhospitable than imagined by many, a reconsideration of planetary protection requirements was launched. A 1978 SSB study identified the need for new criteria because the probability of terrestrial organisms growing on Mars was so low that landers conducting initial exploratory visits to subpolar regions did not require terminal heat sterilization [4]. It was argued that quantitative modeling approaches were subject to gross uncertainties in probability estimates of terrestrial microbe growth at different Martian surface locations. NASA also could not specify how many future landings might be made on Mars. Since both of these factors were model inputs, the resulting predictive probabilities of contamination had severely limited utility.

In light of these issues, NASA reevaluated its planetary protection policy during the early 1980s [4] and proposed a new approach that categorizes missions by type (i.e., flyby, orbiter, lander or sample return) and target (ordered by degree of biological relevance) into five classes. For example, Mars landers and probes without life-detection experiments would be Class IVa, and required to meet bioburden limits of 3×10^5 spores/vehicle and 300 spores/m^2. Missions planned for Mars "special regions" (areas where terrestrial life might have a high probability of propagation), or that carry life-detection instruments, must meet additional criteria [3].

In 1992, an NRC report refined the COSPAR approach by drawing a distinction between Mars missions that would search for life (required to undergo Viking-level sterilization) and those that would not (required only to reduce bioburden to Viking's pre-sterilization level). This distinction was later codified and adopted by COSPAR [9].

In 2005, NASA adopted COSPAR's concept of special regions into its planetary protection policy, defining specific parameters like duration (100 years), maximum spacecraft penetration depth (5 m into the crust), and survival limits for terrestrial life (–15°C temperature, or –20°C including margin; and 0.62 water activity, or 0.5 including margin) [10].

At the turn of the century, after the Galileo mission found Europa to be an ocean world, the SSB was asked to 1) assess the levels of cleanliness and sterilization required to prevent its forward contamination by future orbiters and landers; 2) review methods used to achieve the appropriate level of cleanliness and sterilization for Europa spacecraft, and recommend alternatives in light of recent advancements in science and technology; and 3) identify scientific investigations that should be done to reduce uncertainties of both assessments. The task group concluded that any Europa-bound mission should have a bioload at launch low enough to limit, after the additional

Table 1. 7-step binary-tree algorithm recommended by National Research Council in 2012 for application to OW missions. A "yes" for any question would obviate expensive sterilization measures.

Liquid water	Do current data indicate that the destination lacks liquid water essential for terrestrial life?
Key elements	Do current data indicate that the destination lacks any of the key elements (i.e., carbon, hydrogen, nitrogen, phosphorus, sulfur, potassium, magnesium, calcium, oxygen, and iron) required for terrestrial life?
Physical conditions	Do current data indicate that the physical properties of the target body are incompatible with known extreme conditions for terrestrial life?
Chemical energy	Do current data indicate that the environment lacks an accessible source of chemical energy?
Contacting habitable environments	Do current data indicate that the probability of the spacecraft contacting a habitable environment within 1,000 years is less than 10^{-4}?
Complex nutrients	Do current data indicate that the lack of complex and heterogeneous organic nutrients in aqueous environments will prevent the survival of irradiated and desiccated microbes?
Minimal planetary protection	Do current data indicate that heat treatment of the spacecraft at 60°C for 5 hours will eliminate all physiological groups that can propagate on the target body?

reduction from space irradiation in flight, the final probability of contaminating the Europan ocean to 10^{-4} ([9]). The SSB report recommended that compliance be determined by a "Coleman-Sagan" calculation of contamination risk, which multiplies initial bioload by reduction factors including spacecraft cleaning, exposure to the space environment, and likelihood of introduction into a potential habitat.

In 2012, the NASA Associate Administrator for Science asked the SBB to review and update the 2000 Europa report's recommendations, including application to other icy bodies in the solar system. This request was precipitated partly by Cassini's determination that Enceladus, Titan, and possibly other Saturnian icy moons were ocean worlds also. The board found that reliance on the Coleman-Sagan formulation was not useful because its multiplicative factors may lack statistical independence, and many are prey to uncertain magnitudes. The SSB recommended instead a sequence of seven binary decisions unique to each mission and destination (Table 1), that reflect the geological and environmental conditions of the target body in the context of the metabolic and physiological diversity of terrestrial microorganisms [11]. A 'yes' conclusion for any criterion would release the mission from mitigation activities beyond "routine" cleaning procedures and microbial bioload monitoring. Only if all criteria were determined to be 'no' would a project subject the entire spacecraft to a terminal dry-heat bioload reduction step (heating to >110°C for 30 hours).

The FC issue acquired yet more urgency as multiple habitability-science missions, to diverse destinations, came to be formulated, especially Mars 2020 (planned to cache scientifically selected samples for possible future return to Earth), and Europa Clipper. In 2016, the NASA Science Mission Directorate again asked the SSB to review and assess the policy development process with respect to defined and anticipated needs, and to recommend planetary protection actions. In 2017 the SBB published an interim report detailing the rationales and goals of planetary protection policies, and suggesting a working definition of planetary protection consistent with these rationales. It concluded that the goals of planetary protection, and the rationales supporting them – including protection of Earth, protection of solar system bodies from biological contamination that could compromise scientific investigations of extant or extinct life, and safeguarding of scientific investigations of other solar system bodies – remain mostly unchanged [8].

Throughout the history of space exploration, various risk management models have repeatedly converged on 10^{-4} as the appropriate requirement for limiting forward contamination. Nonetheless, NASA anticipates that "these requirements will be refined in future years" [12]. Such refinements could be informed by evolving understanding of how life functions in extreme environments, knowledge about the target environments, quantification of mission operations concepts, and progress in sterilization technologies for sensitive components and materials.

One point of reference for a multi-faceted community reconsidering its requirements for microbial contamination might be the evolving debate surrounding hospital sterility assurance. A reconsideration of their requirements and policies is currently underway [6], on the basis of studies demonstrating that patient outcomes are not materially affected by three-order-of-magnitude differences in applied sterility assurance levels. Other factors have been found to dominate the experienced infection rate. Since the magnitude of sterilization "overkill" built into typical hospital processes has operational cost implications (e.g., in dollars and time delays to complete sterilization cycles), the community is actively deliberating this risk-management tradespace.

Planetary protection risk management could be subject to a similar phenomenon, in which mission operations factors could dominate the contamination probability far more than spacecraft sterility at launch would. If this is true, sterilization "overkill" alone would burden projects with ineffectual, and avoidable, cost and schedule penalties. Trade-offs should adequately assess and balance the major factors contributing to FC risk, particularly given the continuously evolving understanding of microbial survival and viability, and continuously improving processes for space system integration and verification.

Why Reconsider the Subject Yet Again?

The FC challenge, and how to meet it, comprise a complex subject. Yet today's discussions (most of which occur in technical forums like committee meetings and science conferences) focus almost exclusively on protection-of-science. This purely scientific motive, i.e., avoidance of compromising the study subject by the act of studying it, is only one view. The ethical dimension – what it would mean existentially, philosophically, to alter a non-Earth form of life found somewhere in the cosmos; or alternatively, whether humans should spread Earth life everywhere possible – are largely elided. For example, today's PPOSS project (Planetary Protection for the Outer Solar System) funded by the European Commission's H2020 initiative includes astrobiologists and aerospace professionals, but no ethicists. Intellectual rigor alone might suggest broadening the scope of today's conversation. But two other developments also argue strongly for a broader and continuous reconsideration.

First, we stand at the threshold of the 'century of exploring ocean worlds.' This is not only because we have discovered evidence for so many ocean worlds in the solar system (Figure 1), but also because we either have or can develop the means to explore them, searching throughout their plethora of niches for evidence of life. Doing this satisfactorily could easily take one or two centuries, but not likely more. Once done, humanity would have in its hands the only tangible evidence we will ever have for life elsewhere. Exoplanet spectra may tantalize us and exoplanet images may eventually be obtained, but all of this would still be circumstantial evidence subject to confirmation and explanation only upon physically reaching exoplanet systems in some future too distant and unforeseeable to be relevant.

Returning samples from our solar system's ocean worlds someday would require humanity to be ready, technically and societally, to receive them. Simply put, specific plans for bringing biologically relevant material to Earth will inevitably catalyze widespread societal attention and debate, via the media. A policy of acceptance that purports to adequately protect Earth, but which is written by scientists for scientists, may be inadequate to win that debate. Socializing the FC topic and requirement for missions that do not propose sample return from ocean worlds would allow our technical community to exercise many of the same stakeholders and technologies first, on a problem that carries zero physical risk to humanity.

Second, progress in multiple biological sciences over the past half century has made it clear that life is more diverse and tenacious, and yet more interdependent, than we used to know (Figure 3). Three major research areas appear relevant to the question of alien life.

One of the most exciting frontiers in science today is the exploration of chemistry that is active at the blurry edge of life. One example is retroviruses, encapsulated macromolecule complexes that cannot replicate until becoming incorporated (as proviruses) into the genome of higher organisms. HIV brought retroviruses to

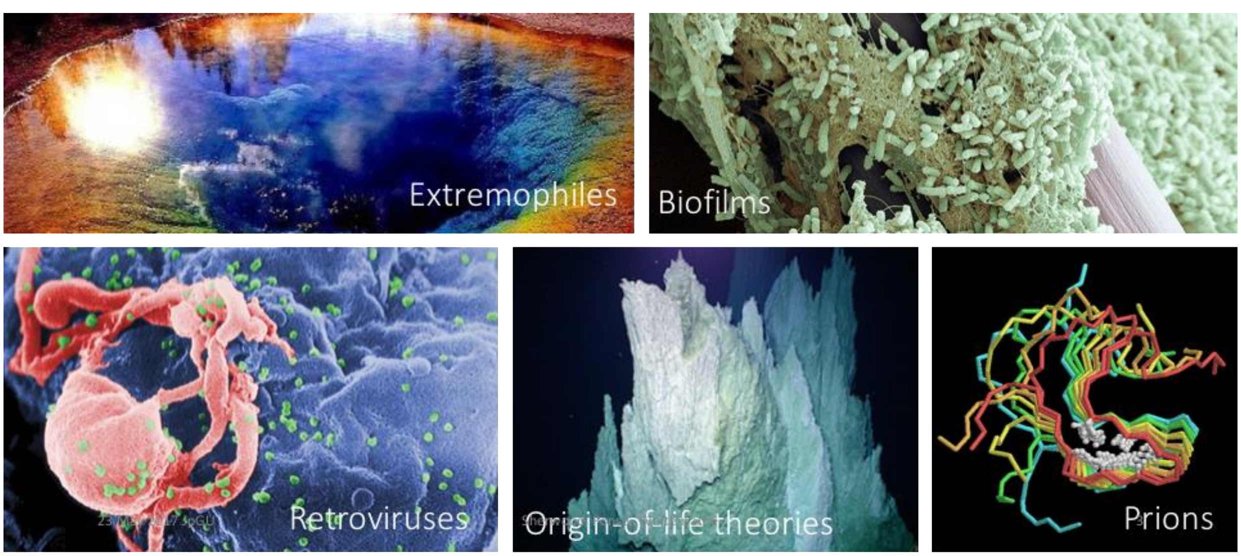

Figure 3. Life is more diverse and tenacious, yet more dependent on ecology, than humanity thought at the time the current forward contamination requirement originated.

widespread attention in the early 1980s. Another example is prions, misfolded configurations of functional proteins, which can proliferate pathologically in living systems, causing devastating dysfunctionality. Creutzveldt Jakob Syndrome and "mad cow disease" brought prions to widespread attention in the 1990s.

Second, extremophiles are now found in, and cultured from, countless niches whose range of conditions far exceeds those found to exist, or to have existed, in several places in space already. On Earth, cells are known to thrive in hot, high-pressure cracks in sub-seafloor rock; deep inside salt and metal mines; in acid hot springs and alkaline seafloor hydrothermal systems; inside radioactive reactor coolant systems; and on equipment exposed to space for years. Microbial cells often demonstrate robust repair mechanisms, and encysted cells can survive otherwise uninhabitable conditions for many years. (And relevant for biological cleanliness standards, diverse biomolecules including DNA are now routinely, forensically extracted from deep-time archeological samples.)

Third, origin-of-life theories have advanced far beyond Darwin's "warm little pond" and Harold-Urey. Contemporary hypotheses debated include wet-dry cycling, and hydrothermal-chimney separation, for ways redox energy might first have been harnessed and active membranes might first have functioned, across the wide range of solar system environments we will have to test in seeking evidence for a second origin.

Already more than halfway through the 21st century's second decade, we can only conclude that, compared to humanity's 1960s, or 1990s, or even 2000s state of knowledge about life – what it is, how it works, where it might be – we have ample reason to be methodical and technically cautious as we venture into the ocean worlds. Maybe Earth life could be more viable elsewhere than we have thought.

Or maybe not; viability might be lower than we have assumed, based on the emerging field of how organisms form, sustain, and depend on communities. Appreciation of ecological relationships, arising first in the 1960s with strong roots in philosophies of ethics, is today a quantitative, multi-dimensional science. Molecular assays of whole communities are a de rigueur analytical tool for studying extreme environments today. Inevitably this is leading to sophisticated questions about how life is managed in real-world settings, or could be.

One direction of research investigates dependencies between microbial life and crystal geology. Another major thrust, with implications throughout our technological civilization, is biofilms, structures built by microbial life to condition a self-conducive environment. Biofilms are tenacious which should motivate caution with respect to forward contamination risk; yet a "single viable organism" introduced into a potential habitat without the protection of a community might not survive long. Much remains to be understood before confident models of Earth organism survivability in alien niches can be validated. It may turn out that our current thinking is far too conservative. It will certainly turn out that what we know at any time will be overcome by new knowledge within just a few years.

We conclude that, without a renewed and continuous conversation that brings all these threads together, for consideration by a broad stakeholder coalition, our technical community cannot systematically, accurately, or even reasonably anticipate the range of manageable futures as a function of potential decisions we might make over time. Not having that strategic analysis then precludes effective preparation, and puts planning at risk.

Analysis of the Forward-Contamination Risk

Wherein lies the actual FC risk?

Risk-management professionals differentiate risk from hazard, because perceptions of risk are so labile as to be an unreliable basis for attaining concurrence on policy. Specific hazards, on the other hand, can be measured and studied. We adopt the definition of Kates & Kasperson [13]:

> *A hazard is "a threat to people and the things they value."*

There is exactly zero threat to people from forward contamination of ocean worlds (except perhaps for the case of Mars, which people may inhabit someday, but even in that case a FC hazard to humans stretches credulity).

That leaves the things we value. By embarking on the exploration of space to seek out life, we must be willing to accept some hazard of violating a thing important to us. What might that be?

First is the protect-the-science objective described above. A societal commitment to avoid destroying or irreversibly complicating future scientific analysis of a potential habitat would largely be a function of societal adherence to the value of scientific principles altogether. So far, the international spacefaring community has a good track record of securing agreement based on science.

Second is a moral obligation, called by Star Trek "The Prime Directive", whose objective is to avoid interfering with living systems or habitats upon first contact. Alas, this branch of "things we value" is not at all amenable to scientific reasoning. It is, however, a great subject for ethics.

We argue here that a modern ethics discussion informed by the most contemporary science would be both advisable and directly beneficial for an OWEP. Advisable, because it would forestall the possibility of a disruptive reexamination or reconsideration far downstream in development of multiple, expensive flight projects; and because it is the right thing to do. Directly beneficial, because it would catalyze international collaboration for program implementation that, in turn, would attract interest, funding, and research opportunities to the participating actors. No 'century of exploring ocean worlds' can occur without such global support.

Sidebar: How does 10^{-4} compare to other small numbers?

Table 2 orders the probabilities of various positive and negative potential events, to establish a mental framework for comparing acceptable risk with desirable return. It particularly helps this comparison for the data to cross fields. For example:

- You have a lower chance of dying this year from an injury than you do of becoming a NASA astronaut, should you apply. For most people, this might mean they should worry less about the risk of dying from injury.
- The probability of NASA contaminating a potential alien habitat, given today's forward-contamination regulation, is about the same as your chance of dying this year from assault with a gun.

Does that seem about right to you? Should our planetary protection policy be stricter than this, or more lenient? What value are you willing to put on protection of potential habitats? In contemplating these questions, you enter the stakeholder conversation that we suggest needs to take place.

Risk Perception and Value Perception

Industrialization has introduced society-level hazards. The field of Risk Management was born of late 20th century circumstances: environmental effects of chemicals and nuclear processes, and their impact on living standards in a democratically vocal society.

From the Risk Management field's many enlightening discoveries, we can derive lessons for considering how to think about the "hazard to things we value" of potential forward contamination of potential alien habitats:

- Low-Probability, High-Consequence risks are assessed by people in a way that is part technical and part psychological and social, or "psycho-social" [14]. This means that their assessments depend on parameters of human perception and evaluation that are in part independent from technical fact or analysis.
- Quantitative risk-assessment tools are inherently limited. If the probability of a low-probability event is exceeded by the uncertainty in the risk model, then decisions should not be based on the analysis.
- People judge very low or very high numbers very poorly. Limitations of human cognitive perception cause distortions in valuations that presume to be based on quantitative comparison. (See *Appendix: Armamentarium Against Human Misjudgment of Quantitative Extremes*)
- Across the field, the dominant driver of perception of risk has turned out to be "distrust of the professional expert, and, by extension, distrust of the process of identifying and dealing with risks" [15]. This can be a devastating factor, able to cripple a field as it did the US nuclear power industry. It is easily understood in the context of post-WWII history. From the 1950s through the mid-1960s, 'white-coated scientists' were a respected priesthood; today, hyperspecialists are viewed far more skeptically: as proffering only one of many valid viewpoints, limited in outlook, and perhaps even under the influence of special interests.

Alerted by the first and second of these lessons, we provide the appendix as 'sensitivity training' for the third, and discuss the fourth below, in the section "Guidance from Precedent."

Framework for a Conversation

In such a parlous decision environment – technical issues that are hard to understand, unlikelihood of accurate modeling, multiplicity of viewpoints including non-technical factors, poorly applicable cognitive tools, and eroded trust in scientific leaders – how could a shared perception and valuation of the hazard be developed, and socialized into policy?

Table 2. Range of probabilities representative of various possible events.

1 in 15	Admission to Yale, 2016 [16]
1 in 20	Lifetime death from injury [17]
1 in 133	Admission to RuPaul's Drag Race, 2017 [17]
1 in 606	Lifetime death from vehicular injury [17]
1 in 1525	Admission to the NASA astronaut class [18]
1 in 1615	Yearly death from an injury [17]
1 in 9737	Lifetime death from aircraft accident [17]
1 in 10,000	**Max allowable, introducing one Earth organism into a potential habitat**
1 in 11,207	Yearly death from assault with a gun [17]
1 in 141,571	Yearly death from falling down stairs [17]
1 in 13,744,732	Yearly death from lightning [17]
1 in 13,983,816	Winning 6-number lottery from pool of 49 numbers [19]

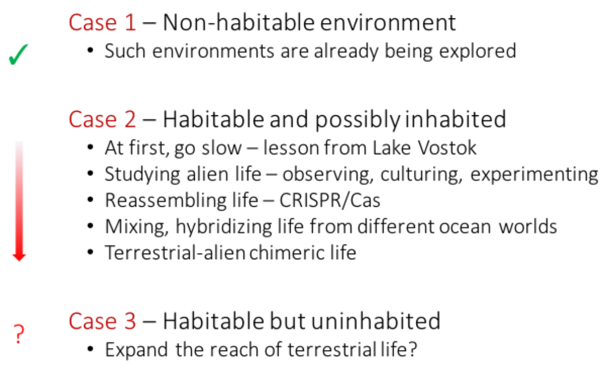

Figure 4. Simple decomposition of forward-contamination cases, combined with current events, reveals the 'sliding scale' of ethical comfort that applies to inevitable decisions.

A first step could be to recognize that value perceptions are themselves labile for this topic; they change along an 'ethical sliding scale' (Figure 4). The figure lays out a scale applicable to FC:

Case 1 captures most scientific space exploration today: we routinely send probes into places we have no basis for thinking might be habitable. For example, even our most extensive exploration of Mars so far has only met with habitable environments paleontologically, by finding evidence in rocks that billions of years ago there was flowing, standing water with habitable chemistry, for a geologically significant duration. Our FC PP process validates that this means there is no present hazard, and evaluates requirements for missions under subcategories of Category IV. Missions into places even less complex are granted less restrictive categories, even more easily. The world does not worry about them.

Case 2 is far more interesting (because it poses ethics problems): places that might be habitable based on our best understanding, or hypothetically, that are confirmed to be inhabited. Most (but not all) people would agree to "go slowly" in exploring such niches, taking care to learn about the environment in increments small enough to avoid irreversible contamination. For Mars, such a sequential protocol is detailed in Sherwood [20].

A sobering example of how easy it is to create scientific confusion is the breeching of Lake Vostok. Earth's largest Antarctic subglacial lake, and overlain by ice deposited over the past 400,000 years, Vostok was isolated from the rest of Earth's biosphere for between five and twenty million years. It was first penetrated in February 2012 by a Russian project; a previously unknown microorganism was discovered. But significant scientific and ethical controversy still attends the ongoing drilling project because of confirmed contamination of the lake water by non-sterile drilling fluids.

How slow is slow enough? How can agreement be reached on standards and the verification of equipment to meet them? What happens when agreement is not reached?

Now suppose that contamination is controlled satisfactorily in exploring an alien habitat, and then life is found, and then confirmed to be native. Then what? First humankind would observe it. Then, we may attempt to culture it, to allow us both to control its growing conditions and expand its study. Ultimately, we may experiment on it directly: e.g., by varying conditions to learn its behaviors and limits. This strategic arc shifts the value we place on the life: from something we watch from a respectful distance, to something we experiment on in our laboratories. This natural evolution – of steps in human-controlled interactions with life forms – inexorably introduces significant ethical implications.

The next natural technical experimental step would be to reengineer the life's chemical machinery. We do this within our own ecology already. The CRISPR/Cas technique is a facile tool for reassembling the chemistry of Earth life, so we would seek to do the same with an alien ecology. Then, as on Earth today with three-parent babies and human DNA spliced into animal models, we might seek to mix chemistries, hybridizing life among the ocean worlds. A final step might be the design of terrestrial-alien chimeras. Where along this scale do you become uncomfortable? Is this a conversation we should have all along the way?

In some ways, the most ethically interesting case is **Case 3**: places that, if found to be habitable, are nonetheless found to be uninhabited. What to do then? Some would say, "Leave it alone." Others would say, "Spread Earth life everywhere we can." What ethical obligations might we perceive in expanding the reach of terrestrial life out into the universe? And how might we reconcile this with our respect for other life unto itself? How would we decide?

At the very least, these are meaty questions, requiring collective action to resolve. Without open discussion, humanity could face the exobiology frontier divided on how to proceed.

The Ethical Dimension

Ethics is a discipline that attempts to apply objective rigor to consideration of moral decisions. The discussion that follows presupposes agreement that there is some moral dimension to decisions that could result in contaminating an alien habitat with Earth life.

A starting point could be a meta-ethical perspective – one that examines different ethical traditions for their potential relevance to the problem. One of the challenges for any ethical treatment of FC, which indeed makes it a rich subject for ethicists, is the fact that the party being subjected to the most potential harm – alien life – falls outside the scope of most ethical theories. Most such theories deal with problems arising in interactions between sentient beings, hence they fall short to address

Sidebar: Why is there an ethical dimension?

A decision becomes an ethical problem when at least two positive values are weighed against each other. If there are no alternatives (i.e., no choice), there can be no ethical problem; and unless at least two values are positive, there can be no ethical problem.

The positive values for FC planetary protection are diffuse:

- The value of research in an uncontaminated pristine environment
- The value to research of an uncontaminated pristine environment
- The intrinsic value of untouched environments
- A possible human obligation toward extraterrestrial life
- The value of minimizing cost and other obstacles to progress

All of these may matter to many stakeholders; but they will matter in 'mixture ratios' as diverse as the stakeholders, and stakeholders will get their information from both the scientific literature and the media. The results of a collaborative ethical exploration cannot be predicted and should not be presumed. But agreeing to a 'recipe' that balances the positive values is essential for PP policy development to be collaborative and broad, so that OW mission development can be efficient and stable.

the party most likely harmed in the case of FC. Today, in advance of exploring potential habitats on faraway worlds, it is simply impossible to meaningfully constrain the characteristics of extraterrestrial life, or even assert anything definitive about the environmental conditions that could enable it. Thus we cannot say anything definitive about the nature of alien life, let alone its sentience or intelligence. Importantly, this means that addressing the ethical dimension from the viewpoint of possible extraterrestrial life may be philosophically interesting, but would be ethically unproductive. This limitation deserves mention because many ethical theories may arise in discussions as if they were arguing "from the outside in." *

The only perspective we have is one centered on the human experience. So every ethical discussion is bound by moral and ethical dimensions that are relevant to, and can be discussed within, human societies. The meta-ethical approach would describe the ethical perspectives that could be applied to those values. Even though the current requirement is a numerical limit on the probability of contamination, quantitative risk assessment is not the only way to approach this ethical problem. What other stances could be taken? And what stakeholders might be behind these ethical positions, and therefore expect a voice in the discussion?

A first step is attempting to ascribe concrete values to the various ethical and moral perspectives on the problem, which requires identifying the social and societal place of these values and the stakeholders associated with them.

The primary value in considering FC has traditionally been protection of science. The act of space exploration destroying latent scientific findings is a well-recognized concern. Scientific integrity as the basis of good science has a communicable societal value. It can be ethically and morally valued within the community of scientists, and also communicated to the broader public and societal institutions, and thus has value there. So scientific integrity is a strong value that has played, and does and will continue to play a role in FC policy-setting. It is hard to give it concrete value, though. For example, progress in detection methodologies, capable of identifying terrestrial genomic material, could lead to a more relaxed approach to FC, but the evolving understanding of extremophile tenacity might tip the balance the other way. Yet this is an entirely "intra-science" discussion best approached within the scientific stakeholder community. Ascribing it to a historically elaborate ethical school would make no sense.

The intrinsic value of an untouched environment is a completely different problem. As this value ascription is done from within human society, ethical theories come into play. Ascribing value to untouched, barren, lifeless places in the universe is a human endeavor that tries to argue for values outside of the scope of human existence. It transcends human existence and thus has to rely on principles that go above social human interaction. This is, for example, where ethical theories that rely on absolute values come into play. One such idea is the notion that the universe was made (created) in its current form and should not be changed. This idea can be found among all major human religious traditions. It is usually framed in a way that compares human agency in spreading life throughout the universe with a deity's work and thus is taboo. While compatible with the idea of preserving pristine research conditions for scientific purposes, its ethical foundation does not stem from societal consensus but rather from revelation or ethical supposition.

One example for this is an argument by philosopher Robert Sparrow [21] against terraforming. In this argument, changing a whole planet to suit humanity's need would demonstrate aesthetic insensitivity and the sin of hubris, both defects of character. Thus, the

* Applying such theories to backward contamination is an even more interesting problem, but outside the scope of this paper. Probabilistic risk assessment is not the only way to address the problem of backward contamination; applying it requires assigning a value to the possible extermination of all life on Earth. That such a valuation for FC is also lacking is evident from comparing the 10^{-4} value to other low-probability events in everyday life (Table 2).

argument relies on some outside value source which is based on the human character. This is aligned with the school of **virtue ethics**, which draws its ethical values from human character. Sparrow argues that the two defects of character basically lead us down the path of trying to become gods (p. 233), an idea that bases ethical values outside of societal consensus or scientific discussion.

This value could also be ethically justified without resorting to transcendental means. The intrinsic value of barren landscapes could be a hypothetical value for potential evolution of life without human interference, in accordance with the ethical principle of Kant's categorical imperative that one should "act only according to that maxim whereby you can, at the same time, will that it should become a universal law." [22] Stated in this way, respecting the current state of untouched environments can be understood as an absolute value in accordance with a **deontological ethics** – an ethical school that bases ethical value ascriptions in duties or rules that must be followed. Under this framework it would be humanity's duty to leave the universe untouched, unsoiled, so as to leave natural processes to their course. This view might be aligned with scientific principles, but only if science restricted itself to pure discovery and observation, without manipulation of its subject matter.

A possible human obligation toward extraterrestrial life is an even more interesting and diverse ethical problem. Several ethical models could be applied to this idea. For example, it could be argued within deontological ethics that life is one of the rules of the universe, and thus needs to be supported and helped along, and that there is therefore a human duty to spread it.

Virtue ethics might argue the same thing from the perspective of the human character. The virtue of being human would necessitate that we, as human beings, help spread life in the universe. What makes this position so interesting for a discussion about forward contamination is the notion that the virtues used to define the value of a certain ethical position are themselves culturally constructed: this position would attempt to represent a diverse stakeholder base that most likely would not all share the same idea of virtues. Griffin [23] argues against a weakening of virtue ethics by only turning inward toward human virtues, by directing the perspective of virtue ethics outward toward the "goods and the ills of the world". This ethical perspective would include the environs of virtues as a basis for debating FC policy, and would be of great interest to a broad range of societal stakeholders.

A **Utilitarian** perspective [24, 25] might argue that in the long run, spreading life and or feeling obligated toward extraterrestrial life would spread happiness (*utilitas*) throughout the universe. This assumes that extraterrestrial life would become sentient, and thus be able to experience, and therefore to produce *utilitas*. The argument, however, is something that has cropped up, and will again, in discussions of human obligations toward alien life. Utilitarianism already plays a role in discussions of the societal value of human space exploration, aimed not at FC but rather at societal needs. [26]

A common religious phrasing of the ethical position of human obligation toward extraterrestrial life is the idea of 'stewardship'. Found in different religious traditions, but also in environmental ethics, this idea will inevitably play a role in ethical discussions about FC [27]. Given its aspect as a religious rule, it might also pose one of the more strict ethical positions within a broad-based discussion.

Ethical values that are aligned with economical values stem from societal consensus, or at least societal interest groups. Such values are the easiest to assess, as they represent economic power and thus societally translatable currency. Aligning them with some of the less socially based ethical systems mentioned above could be a major benefit of the broad stakeholder discussion we propose.

The **economic value view** also has the highest affinity to risk-based ethical assessment, since the notion of economical values can directly assimilate risk into its world-view. This explains the common perception today that the antagonist of planetary protection is mission cost (e.g., budget, schedule, risk to hardware). But reducing the discussion to risk assessment excludes other relevant historical and contemporary positions. Additionally, a key problem for the risk-assessment method of value ascription is the lack of definitive data about the probability of extraterrestrial life and its properties, or even the probability of terrestrial life surviving in extraterrestrial environments. Bachmann & Rippe [28] discuss similar examples and the effect of lack of information on ethical decision making in the field of risk-ethics. While the methodology of risk assessment is familiar, lack of data regarding FC makes ethical evaluations based on risk assessment quite problematic to communicate socially. These challenges are elaborated in the next section.

This brief meta-ethical overview of perspectives that might arise from various ethical positions, perceptions, and value ascriptions reveals the wide range of possible stakeholder orientations, some of which may bring deeply held viewpoints about the implications of FC. Whereas planetary protection discussions to date, even those only about FC, have been predominantly technical and grounded in quantitative risk assessment, it is reasonable to expect that non-technical stakeholders' views will seek a voice.

Reconciling diverse viewpoints to achieve a consensus basis for agency policy-setting cannot be

Table 3. Value clusters act as seeds for the formation of social alliances in an ethics-based decision process.

Care Prime Directive Studying a pristine environment	⇔	Public mission advocacy Mission approval Low Cost
Scientific inquiry Caution Adaptable rules	⇔	Exploratory curiosity Lack of delay Verifiable requirements
Technological advancement	⇔	Less Complexity High assurance of success
Cosmic diversity	⇔	Spreading Earth life
Ideology	⇔	Ideology

straightforward: the values that drive an ethical decision-making process do not line up in a simple matrix that technical managers could use to argue one group of stakeholder interests against another. Rather, the values that drive the ethical discussion form clusters (Table 3). By attracting the interests of certain societal groups and institutions, these value clusters can help form social alliances.

The first group of clustered values all argue for strong FC prevention measures: care for extraterrestrial environments, a "Prime Directive" doctrine of non-interference, and scientific interest in researching a pristine environment without the complication of contaminants confounding measurements. In an ethical decision-making process, this value cluster would be in tension with a cluster capturing institutional and public interests in keeping missions simple and cheap enough to get approved in the first place, which may be related to sustaining public interest and advocacy. While individual values in this tensioned framework might be held by different societal stakeholders, they may equally well be represented within the same group of stakeholders. For example, both the space mission community and the public are, all at the same time, interested in public support, cost containment, scientific integrity, and care for potential habitats. It is precisely this opposition of positive values that makes FC an ethical problem as much as a technical one (see sidebar).

The blending of stakeholder interests is even more pronounced in the second group of value clusters. Here, the values on both sides all represent sound scientific and engineering practice, and the tension between the clusters has characterized science and engineering problems for many decades. But they also may crystallize and represent potential involvement in space missions by other societal groups. For example, arguing sheer exploratory curiosity versus rigorous scientific inquiry characterizes a newly relevant tension between anxious investors in space exploration and traditional government-sponsored scientific exploration. Stakeholder debate between the science and exploration value clusters is already occurring with respect to potential use of space resources and Mars colonization [29].

While the first three rows of tensioned, clustered values represent somewhat familiar debates within science and engineering communities, and between them and other societal stakeholders, the last two rows broaden the frame. The ethical side of FC may render down to a tension between stark, cosmically significant positive values: promoting cosmic diversity versus spreading Earth life. Aligning this simplification with the tensioned clusters above hints at a problem lurking behind the whole discussion, one that every scientific community faces repeatedly: Do we face an ideology problem? Is the decision we try to make something that is clearly in line with our basic principles, or are we influenced by opposing ideologies? The community motivated to have workable planetary protection policies must tackle this challenge. Only by conversing with a broad spectrum of societal stakeholders can ideological problems be resolved or overcome. How does a community learn its own context, and boundaries circumscribing its own decision making process, if not from the outside?

Race [30] has been making this point when considering decision making processes in the event that extraterrestrial life is discovered. Her description that "it is important to recognize that current deliberations and decision making are almost exclusively in the realm of scientific and spacefaring elites" makes it clear that ethical debate must become "planetary in scope". An ethics concerned with space exploration must rely on the equivalent of "informed societal consent". Such consent necessarily relies on open and public discussion including ethical perspectives that encompass a broad spectrum of societal interests.

Guidance from Precedent

The broader the conversation becomes (i.e., the more it includes non-expert participants), the more critical it is to seek lessons from risk-management practice. Two key lessons learned are: 1) trusting the message hinges on trusting the messenger; 2) trust is hard to gain but easy to lose [14, p. 122]. From analyzing nine case studies, Kammen et al. [15] conclude that uncontrollability and dread, potent as they are in subjective risk assessment, are actually overwhelmed by "distrust of the professional expert, and, by extension, distrust of the process of identifying and dealing with risks." This means that continuous communication is necessary to avoid or mitigate accusations of manipulative secrecy. Trust failure events are especially hard to recover from. Non-expert stakeholders can seize on such events, and use them to reinforce refractory antagonistic biases.

Precedents do exist for resolving new, complex societal risks [31], but the most major examples are not fully apt as templates for how to sponsor or hold a broad stakeholder conversation about the technical and ethical sides of forward contamination.

Secret Approach – As 1950s technology progressed from atomic bombs toward the thermonuclear bomb, a theory arose that the first thermonuclear explosion could ignite Earth's atmosphere, incinerating all life. This gave rise to a risk that had to be evaluated, and accommodated somehow into the decision to conduct the first test, or not. Resolution of the debate was led by the scientific community, albeit completely within a secret world populated only by highly educated people, who determined for the rest of us what actions to take about a hazard unknown to us at the time. This model cannot be applied to FC.

Conversant Approach – More recently, twice before high-energy particle accelerators were turned on, it was theorized that doing so could create a microscopic black hole, which would then grow to consume the entire Earth. This gave rise to a risk that had to be evaluated, and accommodated somehow into the decision to turn on the machine, or not. The first time was in 1999, for the Relativistic Heavy Ion Collider at Brookhaven National Laboratory. The second was in 2008, for the much more powerful Large Hadron Collider at CERN – it was thought the 7 TeV proton collisions in the LHC might create a black hole even though the RHIC's collisions had not.

The conversant approach is socially open, with ample opportunity for opponents to articulate their concerns, and may be litigated in the end. Both particle accelerators proceeded; while the go-forward decisions surprised no one, they were 'topically socialized' first.

> *"…everyone on earth had an interest in the contemplated cosmic hazard. In all these three situations the scientific community acted ahead of the rest of the body politic in raising the issue and by openly setting up a review committee encouraging imaginative thought…*
>
> *[In the third case] One of the concerns expressed by the plaintiffs … is that the CERN reviewing committee, while composed of experts was composed of experts interested in seeing the project go ahead. A concern in the opposite direction is that the plaintiffs are merely publicity seekers who have no substance in their claim." [31]*

Even though in all three cases, it never exactly became clear who should have authority to decide on behalf of all humankind, decisions were coordinated, documented, and implemented. This shows that existing institutions may be applied to manage even spectacular societal risks rich in dread risk, and not well understood. Ultimately for FC, this means that risk management should still be feasible in a societally engaged way, even despite contemporary instantaneous and unmoderated technologies for communicating information and misinformation. Managing this conversation to remain civil and be productive, however, would be a substantive challenge.

Kammen et al. [15] offer lessons that could be applied to a conversation about forward contamination: 1) that analysts go beyond mere "accurate" assessments, considering also the social and economic context; 2) that peer review "be integral to the evaluation of claims of risk;" 3) that strengths and weaknesses of the analyses be assessed using an interdisciplinary approach, and clearly communicated; 4) that interactions between "experts" and the "concerned public" be cooperative, not antagonistic; and 5) that opinions and negotiating positions be kept fluid.

Moving Forward

Elements of a conversation approach would naturally include two major parts:

1. **Information campaign** that socializes the current state and future options (i.e., potential decisions that will need to be made)
   - The types of exploration we can now undertake
   - Implications of those missions for science and for potential alien life
   - How we manage the risk today.
2. **Open, inclusive international dialogue with wide stakeholder reach** that allows evolution and codification of a deterministic policy useful for planners and projects
   - Scientists + Ethicists + Managers + Citizens
   - Thought leaders from all generations.

Table 4 arrays nine types of institutional player involved in public risk management today (after Hood et al. [32], p.141). Examples illustrate the diversity, span of authority, and interests that each type of player might bring to risk-management conversations. So far, international development of the FC requirement has involved four of the types: executives implementing policies determined by legislative bodies, at national and supranational levels. For example, NASA's Planetary Protection Office assures that projects funded by US taxpayer resources comply with today's 10^{-4}

Table 4. Nine types of institutional player (with examples in black font) participate in public risk management. Four players (blue font) are engaged in setting planetary protection policy today. Expanded FC conversation could introduce new players (typified by examples in gray font).

	Institutional Type		
Scope of Authority	**Core Executive Bodies**	**Independent Public Bodies**	**Private or Independent Bodies**
Supranational	European Commission United Nations	US Supreme Court COPUOS, Planetary Protection Panel (PPP)	Greenpeace ---
National	British Parliament NASA, FAA, ESA, JAXA	US Environmental Protection Agency US Senate	National Association of Insurers (NASA Watch?)
Subnational	California state government ---	Los Angeles Dept. of Water and Power ---	Local businesses and policy activists (SpaceX?)

requirement. As a US federal agency, NASA is bound by the Outer Space Treaty, which was ratified by the US Senate. The treaty terms were negotiated under the aegis of the UN COPUOS (Committee on the Peaceful Uses of Outer Space), which adopted the requirement written by its standing Planetary Protection Panel (PPP). A requirement change proposed by the PPP would be put up for COPUOS vote, upon which it would become enforceable under the Outer Space Treaty. ESA and other space agencies operate the same way. No subnational institutions are involved.

However, more institutional types may become involved in the future; examples are suggested parenthetically in the table. For example, SpaceX (a subnational private company) has widely publicized plans to send humans to Mars, an activity with direct implications for forward contamination. Clearly they would seek a voice in conversations about benefit and risk of the PP hazard, and the cost of meeting requirements. Because SpaceX is incorporated in the state of California, this might fold in the state government (subnational executive body) as well. And if NASA begins planning to enter special regions at Mars or other ocean worlds, it is easy to foresee private "watchdog" players (e.g., NASA Watch) also entering the conversation.

No resources are allocated today for holding such a complex conversation, despite the factors described in the introduction: 1) a latent campaign to explore many ocean worlds; 2) a driving requirement with deep roots; 3) a rapidly-changing appreciation for what life is and does; and 4) a rapidly-changing social milieu, whose viewpoints cannot be assumed but who will inherit the consequences of near-term decisions.

Conclusion

We advocate a fresh, updated and ongoing, broad stakeholder conversation about the forward-contamination hazard and how to manage it. Now is a perfect time to either reinforce or reach a new consensus about the assumptions, issues, and driving requirements needed to enable a century's worth of potential ocean world exploration.

This conversation should include all of the stakeholder communities most likely to insist on a voice eventually. It should include two parts: one that informs accurately, successfully, and widely; and one that facilitates dialogue, leading to concurrence and codification in national and international law. It would have to be designed and conducted to provide ample time, potential for interaction, and organizational resources to examine both technical and ethical sides of the issue openly.

The COPUOS PPP meets again next either before or at the next COSPAR Assembly, in September 2018. With the framework outlined here, COPUOS could catalyze a contemporary stakeholder conversation about how to competently and responsibly manage the scientific exploration we are now verging on. The questions raised in this analysis need to be pondered.

A thoughtful, open, and inclusive approach could enable a coherent and proactive ocean worlds exploration program; the contrapositive is equally true.

Appendix: Armamentarium against Human Misjudgment of Quantitative Extremes

Even if managing societal risks was only about making rational decisions based on technical analysis, human comprehension of probability, consequence, and value are subjective and labile; and human decision-making is dominated by heuristics anyway.

Human cognition about LPHC (low probability, high consequence) events in particular is prey to many kinds of errors documented in the risk literature. Taken together, these distortions render classic decision theory useless for explaining human interpretation and decision-making in these situations. "Biases in probability judgment are violations of almost every theory of choice." [33]

Psychometric experiments have yielded predictive heuristics about many faulty cognitive patterns [14]:

- *Identity-line distortion.* The identity line refers to the line with unit slope on a graph of estimated vs. actual frequency of HC events. People systematically over-estimate the frequency of very low-frequency events (e.g., death by tornado in the US, ~10^2/yr in actuality), yet under-estimate the frequency of very high-frequency events (e.g., death from stroke in the US, >10^5/yr in actuality). Estimates in both cases are systematically about an order of magnitude wrong, lowering the slope of the identity line.
- *Conjunction fallacy.* People judge event probabilities by their plausibility. Because detailing an event increases perceptions of its plausibility, it also increases casual estimates of its probability. Quantitative analysis does the reverse, by detailing sub-events whose conjunction is required for the resultant event to occur. "Scenarios rich in detail often have a plausibility that outweighs their likelihood."
- *Optimism.* "It can't happen to me; it hasn't happened to me." Almost 90% of drivers believe they are better than average.
- *Availability.* Media attention leads people to radically misjudge relative probabilities, e.g., the probability of death from heart disease (reported via statistics) vs. death from airplane crashes (reported with lurid video). Even professionals' judgment of the relative probability of branches of a fault tree varies widely as a function of the availability of details about each branch.
- *Over-confidence.* When people are asked to assign confidence intervals to their own estimates, their 90%-confidence interval contains the true value only about half the time. This exacerbates the tendency to report and discuss point-estimates rather than genuine intervals. The only known mitigation is to report only intervals, without any point-estimates.
- *Ignoring low-probability risks.* People have trouble mentally processing extreme numbers. Researchers found a four-fold increase in willingness to wear seat belts when risk was cast as lifetime risk of fatality (0.01) rather than per-trip risk (0.00000025).

Indeed, even perceptions about LP numbers are labile. The phrase "one in a million" was originally used to mean "impossible." One of its first documented uses, in a 1959 conversation between Pearl S. Buck and Arthur H. Compton, was to characterize the expected impossibility of an intentional thermonuclear blast accidentally detonating the Earth's atmosphere [15]. Yet less than a half-century later, we live in a world where a million does not seem like so much.

Kammen et al. [15] summarize other sources of unpredictability and variability in perceptions of risk. All involve departures from classic expected utility theory, in which the integral of utilities of possible outcomes, weighted by their respective probabilities, yields a rational-actor decision. Unsurprisingly, this theory has been found to not model well how people actually make decisions:

- Risk editing. People simplify risks so they can comprehend them.
- Magnification of imminent events. "People are impatient about the near future and myopic about the distant future."
- Contingent weighting. Heuristic decision-making often assigns weight to outcomes based on their impact, and even to perceptions of likelihood based on their desirability.
- Lexicographic choice. An extreme heuristic, common because it is far simpler for most people, chooses one most-important factor to drive a decision.
- Mental accounting. Tradeoffs across different kinds of mental accounts – e.g., dollars for lives – are especially problematic. Even assessments of personal lethality are labile: for example, identifiable lives count more than equivalent statistical lives.
- Loss aversion. Potential losses count more than potential gains of equal size.

Setting policy for managing technical risks is a psychosocial challenge, not just a technical problem. The litanies of cognitive defects and heuristic patterns described here mean that policies for managing the risk of forward contamination can be neither straightforward, nor necessarily definitive, nor fixed through time.

List of references

[1] B. Sherwood, J. Lunine, C. Sotin, T. Cwik, and F. Naderi, "Program options to explore ocean worlds," presented at the Global Space Exploration Conference (GLEX 2017), Beijing, China, 2017.

[2] B. Sherwood, "Strategic map for exploring the ocean-world Enceladus," *Acta Astronautica,* vol. 126, pp. 52-58, 2016.

[3] NASA, "NASA Procedural Requirements (NPR) 8020.12D," ed. Washington, DC, 2010.

[4] M. Meltzer, N. H. Office, and J. D. Rummel, *When Biospheres Collide: A History of Nasa's Planetary Protection Programs (Nasa History Publication Sp-2011-4234)*: Books Express Publishing, 2011.

[5] L. Jaffe, "Sterilization of unmanned planetary and lunar space vehicles-An engineering examination," Jet Propulsion Laboratory, Pasadena, California1963.

[6] S. W. Srun, B. J. Nissen, T. D. Bryans, and M. Bonjean, "Medical device SALs and surgical site infections: a mathematical model," *Biomedical Instrumentation & Technology,* vol. 46, pp. 230-237, 2012.

[7] National Research Council, "A Review of Space Research," National Academy of Sciences, Washington, DC, 1962.

[8] National Academies of Sciences, Engineering, and Medicine, *The Goals, Rationales, and Definition of Planetary Protection: Interim Report*. Washington, DC: The National Academies Press, 2017.

[9] National Research Council, *Preventing the Forward Contamination of Europa*. Washington, DC: The National Academies Press, 2000.

[10] D. Beaty, K. Buxbaum, M. Meyer, N. Barlow, W. Boynton, B. Clark, J. Deming, P. T. Doran, K. Edgett, S. Hancock, J. Head, M. Hecht, V. Hipkin, T. Kieft, R. Mancinelli, E. McDonald, C. McKay, M. Mellon, H. Newsom, G. Ori, D. Paige, A. C. Schuerger, M. Sogin, J. A. Spry, A. Steele, K. Tanaka, and M. Voytek, "Findings of the Mars special regions science analysis group," *Astrobiology*, vol. 6, pp. 677-732, Oct 2006.

[11] National Research Council, *Assessment of Planetary Protection Requirements for Spacecraft Missions to Icy Solar System Bodies*. Washington, DC: The National Academies Press, 2012.

[12] C. Conley, "Planetary Protection for Icy Moons: updating the SSB Europa Report " NASA, Washington, DC2012. http://www.lpi.usra.edu/pss/may2012/presentations/4Conley_PlanetaryProtection.pdf

[13] R. W. Kates and J. X. Kasperson, "Comparative risk analysis of technological hazards (a review)," *Proceedings of the National Academy of Sciences*, vol. 80, pp. 7027-7038, 1983.

[14] N. Pidgeon, C. Hood, D. Jones, B. Turner, and R. Gibson, "Risk perception (Chapter 5)," in *Risk: Analysis, perception and management*, ed: Royal Society, 1992, pp. 89-134.

[15] D. Kammen, A. Schlyakhter, and R. Wilson, "What is the risk of the impossible?," *Technology: Journal of the Franklin Institute* vol. 331A, pp. 97-116, 1994.

[16] A. Jackson, "Ivy League admission letters just went out — here are the acceptance rates for the class of 2021," in *Business Insider*, ed, 2017.

[17] Insurance Information Institute. (2014). *Mortality Risk*. Available: http://www.iii.org/fact-statistic/mortality-risk

[18] NASA. (2016). *Record Number of Americans Apply to #BeAnAstronaut at NASA (RELEASE 16-018)*. Available: https://www.nasa.gov/press-release/record-number-of-americans-apply-to-beanastronaut-at-nasa

[19] Wonderopolis. (2017). *Wonder of the Day #335: What are your odds of winning the lottery?* . Available: https://wonderopolis.org/wonder/what-are-your-odds-of-winning-the-lottery

[20] B. Sherwood, "Progressive Protocol for Planetary Protection During Joint Human and Robotic Exploration of Mars," in *International Astronautical Congress, Paper IAC-04-IAA*, 2004.

[21] R. Sparrow, "The ethics of terraforming," *Environmental Ethics*, vol. 21, pp. 227-245, 1999.

[22] I. Kant and L. W. Beck, *Foundations of the Metaphysics of Moral: And what is Enlightenment?*: Library of the Liberal Arts, 1959.

[23] J. Griffin, "Virtue Ethics and Environs," *Social Philosophy and Policy*, vol. 15, pp. 56-70, 1998.

[24] J. Bentham, *An Introduction to the Principles of Morals and Legislation*: Dover Publications Inc., 2009.

[25] J. S. Mill, *Utilitarianism (Crisp, Roger, ed.)*: Oxford University Press, 1998.

[26] S. Lingner, "Human spaceflight as a matter of culture and national vision," in *Humans in Outer Space — Interdisciplinary Odysseys*, L. Codignola, K.-U. Schrogl, A. Lukaszczyk, and N. Peter, Eds., ed Vienna: Springer Vienna, 2009, pp. 175-181. https://doi.org/10.1007/978-3-211-87465-3_22

[27] M. Waltemathe, "Probleme, die wohl der Mensch zuvor bedacht hat. Ethische und theologische Aspekte bemannter Raumfahrt," in *Öffentlicher Raum. Theologische, religionswissenschaftliche und ethisch-normative Dimensionen* C. Wustmans, Ed., ed: Sozialethische Materialien, SEM 4, 2016, pp. 189-199.

[28] A. Bachmann and K. P. Rippe, *Ethische Risikobewertung: Ausarbeitung von Kriterien und Instrumenten für eine ethische Risikobewertung im Zusammenhang mit dem Einsatz, insbesondere der Freisetzung von GVO* BAFU-Forschungsprogramm, 2008.

[29] P. T. Metzger, "Space development and space science together, an historic opportunity," *Space Policy*, vol. 37, pp. 77-91, 2016.

[30] M. S. Race, "Space Exploration and Searches for Extraterrestrial Life: Decision Making and Societal Issues," in *Encountering Life in the Universe*, C. Impey, A. H. Spitz, and W. Stoeger, Eds., ed Tucson, Arizona: University of Arizona Press, 2013, pp. 141-156.

[31] R. Wilson, "The development of risk analysis: A personal perspective," *Risk Analysis*, vol. 32, pp. 2010-2019, 2012.

[32] C. C. Hood, D. K. C. Jones, N. F. Pidgeon, B. A. Turner, and R. Gibson, "Risk management (Chapter 6)," in *Risk: Analysis, perception and management*, ed: Royal Society, 1992, pp. 89-134.

[33] C. F. Camerer and H. Kunreuther, "Decision Processes for Low Probability Events: Policy Implications," *Journal of Policy Analysis and Management* vol. 8, pp. 565-592, 1989.

Acknowledgments and Disclaimers

This pre-decisional analysis was conducted by the Jet Propulsion Laboratory/Caltech, under contract to the National Aeronautics and Space Administration.

The cost information contained in this document is of a budgetary and planning nature and is intended for informational purposes only. It does not constitute a commitment on the part of JPL and/or Caltech.